US006590861B1

(12) **United States Patent**
Vepa et al.

(10) **Patent No.: US 6,590,861 B1**
(45) **Date of Patent: Jul. 8, 2003**

(54) **COMBINING VIRTUAL LOCAL AREA NETWORKS AND LOAD BALANCING WITH FAULT TOLERANCE IN A HIGH PERFORMANCE PROTOCOL**

(75) Inventors: **RamKrishna Vepa**, Danville, CA (US); **Roman Baker**, San Jose, CA (US); **Sameer Nanda**, Sunnyvale, CA (US); **James Binder**, San Jose, CA (US); **Thomas Albert Maufer**, Santa Clara, CA (US)

(73) Assignee: **3Com Corporation**, Santa Clara, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **09/272,083**

(22) Filed: **Mar. 18, 1999**

(51) **Int. Cl.**[7] **H04J 1/16**

(52) **U.S. Cl.** **370/216**; 370/218; 370/230; 370/242; 709/238

(58) **Field of Search** 370/471, 401, 370/402, 216, 217, 218, 241, 242, 244, 245, 250, 392, 463, 466; 709/223, 224, 238, 105, 235, 250

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,394,402	A *	2/1995	Ross	370/402
5,864,535	A *	1/1999	Basilico	370/392
5,918,021	A *	6/1999	Aditya	709/235
6,049,528	A *	4/2000	Hendel et al.	370/235
6,052,733	A *	4/2000	Mahalingam et al.	709/235
6,151,297	A *	11/2000	Congdon et al.	370/216

OTHER PUBLICATIONS

Sarparanta, Fault Tolerance—Requirements and Solutions, Helsinki Unversity of Technology, pp. 1–12, May 3, 1998.*
Decisys, The Viertual LAN Technology Report, pp. 1–20, May 1996.*
Liu, Support for 802.1Q/802.1p in the Catalyst 5500/5000 System, Cisco Systems, Inc., pp. 1–17, Sep. 1998.*

* cited by examiner

Primary Examiner—Seema S. Rao
Assistant Examiner—Frank Duong
(74) *Attorney, Agent, or Firm*—Wagner Murabito & Hao LLP

(57) **ABSTRACT**

A method to send and receive data packets over a network of computer systems (e.g., a plurality of virtual local area networks [VLANs] compliant with the IEEE 802.1Q standard) using a server computer system adapted to have a plurality of network interface cards (NICs) coupled thereto. An outgoing data packet is addressed using a first media access control (MAC) address, wherein the first MAC address is a virtual MAC address representing the plurality of NICs. A NIC is selected from the plurality of NICs using a load balancing scheme. The functionality of the selected NIC is verified using a fault tolerance scheme. The first MAC address in the outgoing data packet is replaced with a second MAC address, wherein the second MAC address represents the selected NIC. The outgoing data packet is then transmitted using the selected NIC. For an incoming data packet received at one of the plurality of NICs, the existing MAC address in the incoming data packet is replaced with the first MAC address. Thus, the present embodiment of the present invention provides a high performance method that supports multiple VLANs and elegantly combines a load balance scheme and a fault tolerance scheme into a software element implemented on a server computer system.

27 Claims, 10 Drawing Sheets

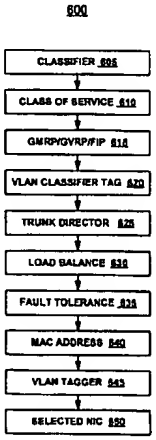

US006560229B1

(12) **United States Patent**
Kadambi et al.

(10) **Patent No.: US 6,560,229 B1**
(45) **Date of Patent: May 6, 2003**

(54) **NETWORK SWITCHING ARCHITECTURE WITH MULTIPLE TABLE SYNCHRONIZATION, AND FORWARDING OF BOTH IP AND IPX PACKETS**

(75) Inventors: **Shiri Kadambi**, Los Altos, CA (US); **Shekhar Ambe**, San Jose, CA (US)

(73) Assignee: **Broadcom Corporation**, Irvine, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **09/343,410**

(22) Filed: **Jun. 30, 1999**

Related U.S. Application Data

(60) Provisional application No. 60/095,972, filed on Aug. 10, 1998, and provisional application No. 60/092,220, filed on Jul. 8, 1998.

(51) **Int. Cl.**[7] **H04L 12/28**; H04L 12/54

(52) **U.S. Cl.** **370/392**; 370/396; 370/428

(58) **Field of Search** 709/215, 224, 709/238, 242, 245; 707/104; 710/52; 370/400, 402, 392, 257, 401, 465, 466, 467, 396, 389, 350, 503, 506, 507, 509–510

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,434,863	A	7/1995	Onishi et al.
5,555,398	A	9/1996	Raman
5,644,784	A	7/1997	Peek

(List continued on next page.)

FOREIGN PATENT DOCUMENTS

EP	0 312 917	A	4/1989
EP	0 465 090	A	1/1992
EP	0 854 606	A2	7/1998

(List continued on next page.)

OTHER PUBLICATIONS

"Catalyst 8500 CSR Architecture" White Paper, Cisco Systems Online, Jul. 1998.

"Computer Networks", A.S. Tanenbaum, Prentice–Hall Int, USA XP002147300 pp. 309–320, 1988.

Primary Examiner—Kwang Bin Yao
Assistant Examiner—Hanh Nguyen
(74) *Attorney, Agent, or Firm*—Squire, Sanders & Dempsey L.L.P.

(57) **ABSTRACT**

A network switch for network communications includes a first data port interface. The first data port interface supports a plurality of data ports transmitting and receiving data at a first data rate. A second data port interface is provided; the second data port interface supports a plurality of data ports transmitting and receiving data at a second data rate. A CPU interface is provided, with the CPU interface configured to communicate with a CPU. An internal memory is provided, and communicates with the first data port interface and the at least one second data port interface. A memory management unit is provided, and includes an external memory interface for communicating data from at least one of the first data port interface and the second data port interface and an external memory. A communication channel is provided, with the communication channel communicating data and messaging information between the first data port interface, the second data port interface, the internal memory, and the memory management unit. A plurality of semiconductor-implemented lookup tables are provided, with the lookup tables including an address resolution lookup/layer three lookup, rules tables, and VLAN tables. One of the data port interfaces is configured to update the address resolution table based on newly learned layer to addresses. An update to an address table associated with an initial data port interface of the first and second data port interfaces results in the initial data port interface sending a synchronization signal to the other address resolution tables in the network switch. As a result, all address resolution tables on the network switch are synchronized on a per entry basis.

34 Claims, 19 Drawing Sheets

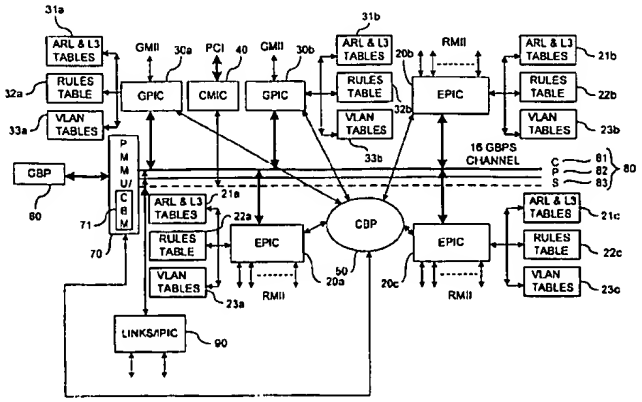

US006104700A

United States Patent [19]

Haddock et al.

[11] Patent Number: 6,104,700

[45] Date of Patent: Aug. 15, 2000

[54] **POLICY BASED QUALITY OF SERVICE**

[75] Inventors: **Stephen R. Haddock**, Los Gatos; **Justin N. Chueh**, Palo Alto; **Shehzad T. Merchant**, Mountain View; **Andrew H. Smith**, Palo Alto; **Michael Yip**, Sunnyvale, all of Calif.

[73] Assignee: **Extreme Networks**, Cupertino, Calif.

[21] Appl. No.: **09/018,103**

[22] Filed: **Feb. 3, 1998**

Related U.S. Application Data

[60] Provisional application No. 60/057,371, Aug. 29, 1997.

[51] **Int. Cl.**[7] **H04L 12/56**

[52] **U.S. Cl.** **370/235**; 370/412; 370/429

[58] **Field of Search** 370/412, 413, 370/414, 415, 416, 417, 418, 410, 229, 230, 231, 232, 233, 234, 235, 236, 237, 401, 402, 428, 429, 445

[56] **References Cited**

U.S. PATENT DOCUMENTS

5,381,413	1/1995	Tobagi et al. .	
5,394,402	2/1995	Ross	370/94.1
5,499,238	3/1996	Shon	370/60.1
5,570,360	10/1996	Klausmeier et al.	370/412
5,629,937	5/1997	Hayter et al.	370/412
5,742,772	4/1998	Sreenan	395/200.56
5,748,614	5/1998	Wallmeier	370/412
5,796,719	8/1998	Peris et al.	370/418
5,850,399	12/1998	Ganmukhi et al.	370/412
5,859,835	1/1999	Varma et al.	370/412

OTHER PUBLICATIONS

International Search Report, PCT/US98/17817, 4 pages.

Primary Examiner—Huy D. Vu

Attorney, Agent, or Firm—Blakely, Sokoloff, Taylor & Zafman

[57] **ABSTRACT**

A flexible, policy-based, mechanism for managing, monitoring, and prioritizing traffic within a network and allocating bandwidth to achieve true quality of service (QoS) is provided. According to one aspect of the present invention, a method is provided for managing bandwidth allocation in a network that employs a non-deterministic access protocol, such as an Ethernet network. A packet forwarding device receives information indicative of a set of traffic groups, such as: a MAC address, or IEEE 802.1p priority indicator or 802.1Q frame tag, if the QoS policy is based upon individual station applications; or a physical port if the QoS policy is based purely upon topology. The packet forwarding device additionally receives bandwidth parameters corresponding to the traffic groups. After receiving a packet associated with one of the traffic groups on a first port, the packet forwarding device schedules the packet for transmission from a second port based upon bandwidth parameters corresponding to the traffic group with which the packet is associated. According to another aspect of the present invention, a method is provided for managing bandwidth allocation in a packet forwarding device. The packet forwarding device receives information indicative of a set of traffic groups. The packet forwarding device additionally receives information defining a QoS policy for the traffic groups. After a packet is received by the packet forwarding device, a traffic group with which the packet is associated is identified. Subsequently, rather than relying on an end-to-end signaling protocol for scheduling, the packet is scheduled for transmission based upon the QoS policy for the identified traffic group.

29 Claims, 5 Drawing Sheets

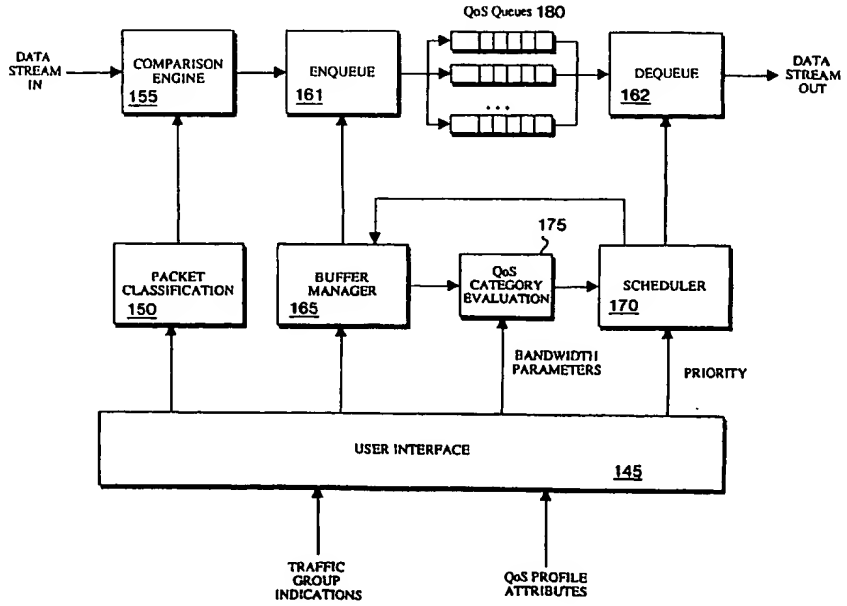

US006104696A

United States Patent [19]

Kadambi et al.

[11] Patent Number: 6,104,696

[45] Date of Patent: Aug. 15, 2000

[54] **METHOD FOR SENDING PACKETS BETWEEN TRUNK PORTS OF NETWORK SWITCHES**

[75] Inventors: **Shiri Kadambi**, Los Altos; **Shekhar Ambe**, San Jose, both of Calif.

[73] Assignee: **Broadcom Corporation**, Irvine, Calif.

[21] Appl. No.: **09/343,718**

[22] Filed: **Jun. 30, 1999**

Related U.S. Application Data

[60] Provisional application No. 60/092,220, Jul. 8, 1998, and provisional application No. 60/095,972, Aug. 10, 1998.

[51] **Int. Cl.**[7] **H04L 12/26**

[52] **U.S. Cl.** **370/218**; 370/221; 370/225; 370/392; 370/397; 370/399

[58] **Field of Search** 370/217, 218, 370/221, 225, 392, 397, 399

[56] **References Cited**

U.S. PATENT DOCUMENTS

5,404,461	4/1995	Olnowich et al.	710/131
5,432,783	7/1995	Ahmed et al. .	
5,838,681	11/1998	Bonomi et al. .	
5,841,771	11/1998	Irwin et al. .	
5,862,136	1/1999	Irwin .	
5,991,297	11/1999	Palnati et al.	370/389
6,016,310	1/2000	Muller et al.	370/255

Primary Examiner—Hassan Kizou
Assistant Examiner—Mitchell Slavitt
Attorney, Agent, or Firm—Arent Fox Kintner Plotkin & Kahn PLLC

[57] **ABSTRACT**

The present invention is directed to a system and method of sending packets between ports on trunked network switches. The method includes providing a first switch having a plurality of communication ports thereupon, and providing a second switch having a plurality of communication ports thereupon. A trunk connection is provided between the first switch and the second switch, with the trunk connection including at least two of the plurality of ports from the first switch being connected to at least two of the plurality of ports of the second switch. A rules table is provided, defining a set of rules identifying which port of the trunk connection will be used for communication. A packet is sent from a first port on the first switch to a second port on the second switch. The packet is received at an ingress submodule of the first switch, and a lookup is performed on one of a source address and a destination address of the packet based upon a lookup table provided in the ingress submodule. It is then identified that the first switch and second switch are connected with the trunk connection by a trunk bit in a lookup entry matched by the destination address. A rules tag bit in the lookup entry is then identified, and the rules tag bit is then compared to a rules table defining which trunk port of the trunk bundle will be used for communication. The rules tag determines which address bits will be used to identify a trunk port for communication. The packet is then forwarded to the destination address on the identified trunk port. The packet is then stored in memory, and then retrieved from memory with an egress unit. The egress unit then forwards the packet to the identified trunk port.

19 Claims, 19 Drawing Sheets

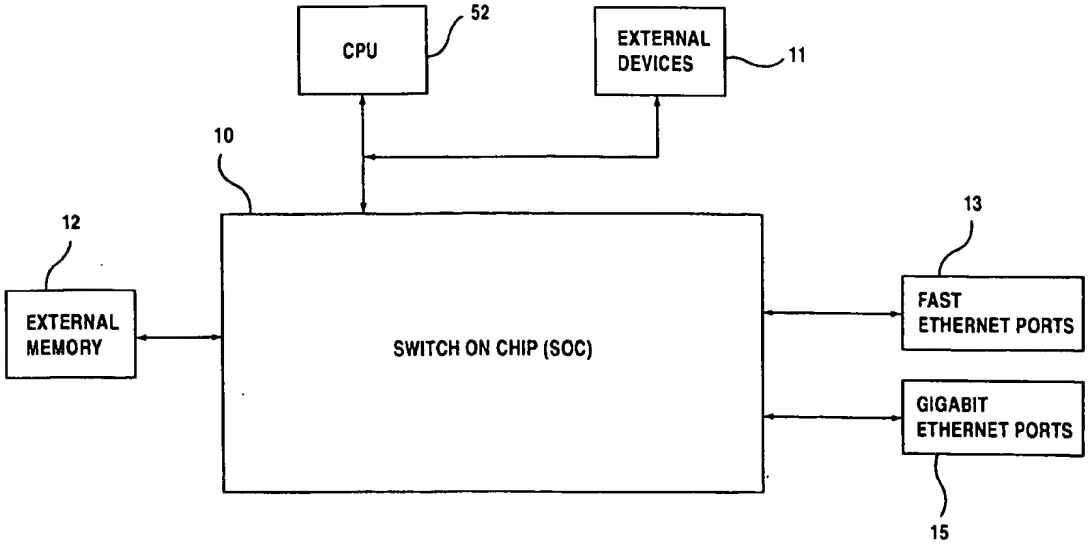

US005909686A

United States Patent [19]

Muller et al.

[11] Patent Number: 5,909,686

[45] Date of Patent: Jun. 1, 1999

[54] **HARDWARE-ASSISTED CENTRAL PROCESSING UNIT ACCESS TO A FORWARDING DATABASE**

[75] Inventors: **Shimon Muller**, Sunnyvale; **Ariel Hendel**, Cupertino; **Louise Yeung**, San Carlos; **Leo Hejza**; **Shree Murthy**, both of Sunnyvale, all of Calif.

[73] Assignee: **Sun Microsystems, Inc.**, Mountain View, Calif.

[21] Appl. No.: **08/885,047**

[22] Filed: **Jun. 30, 1997**

[51] **Int. Cl.**[6] **G06F 17/30**

[52] **U.S. Cl.** **707/104**; 707/1; 370/400; 370/401; 370/402; 370/224

[58] **Field of Search** 707/1, 104; 370/400, 370/401, 402, 229; 395/200.47, 200.48, 200.49

[56] **References Cited**

U.S. PATENT DOCUMENTS

5,309,437	5/1994	Perlman et al.	340/827
5,459,717	10/1995	Mullan et al.	370/351
5,461,624	10/1995	Mazzola	370/402
5,485,455	1/1996	Dobbins et al.	370/255
5,500,860	3/1996	Perlman et al.	370/85.13
5,790,456	8/1998	Dobbins et al.	370/400

OTHER PUBLICATIONS

"Foundry Products", downloaded from Website http://www.foundrynet.com/ on Jun. 19, 1997.

Anthony J. McAuley & Paul Francis, "Fast Routing Table Lookup Using CAMs", IEEE, 1993, pp. 1382–1390.

"Gigabit Ethernet", Network Strategy Report, The Burton Group, v2, May 8, 1997 40 pages.

"IP On Speed", Erica Roberts, Internet–Draft, Data Communications on the Web, Mar. 1997, 12 pages.

"Multilayer Topology", White Paper, Internet–Draft, 13 pages, downloaded from Website http://www.baynetworks.com on Apr. 18, 1997.

Primary Examiner—Paul R. Lintz

Attorney, Agent, or Firm—Blakely Sokoloff Taylor & Zafman

[57] **ABSTRACT**

A method and apparatus for providing hardware-assisted CPU access to a forwarding database is described. According to one aspect of the present invention, a switch fabric provides access to a forwarding database on behalf of a processor. The switch fabric includes a memory access interface configured to arbitrate access to a forwarding database memory. The switch fabric also includes a search engine coupled to the memory access interface and to multiple input ports. The search engine is configured to schedule and perform accesses to the forwarding database memory and to transfer forwarding decisions retrieved therefrom to the input ports. The switch fabric further includes command execution logic that is configured to interface with the processor for performing forwarding database accesses requested by the processor. According to another aspect of the invention one or more commands are provided to implement the following functions: (1) learning a supplied address; (2) reading associated data corresponding to a supplied search key; (3) aging forwarding database entries; (4) invalidating entries; (5) accessing mask data, such as mask data that may be stored in a mask per bit (MPB) content addressable memory (CAM), corresponding to a particular search key; (6) replacing forwarding database entries; and (7) accessing entries in the forwarding database.

23 Claims, 9 Drawing Sheets

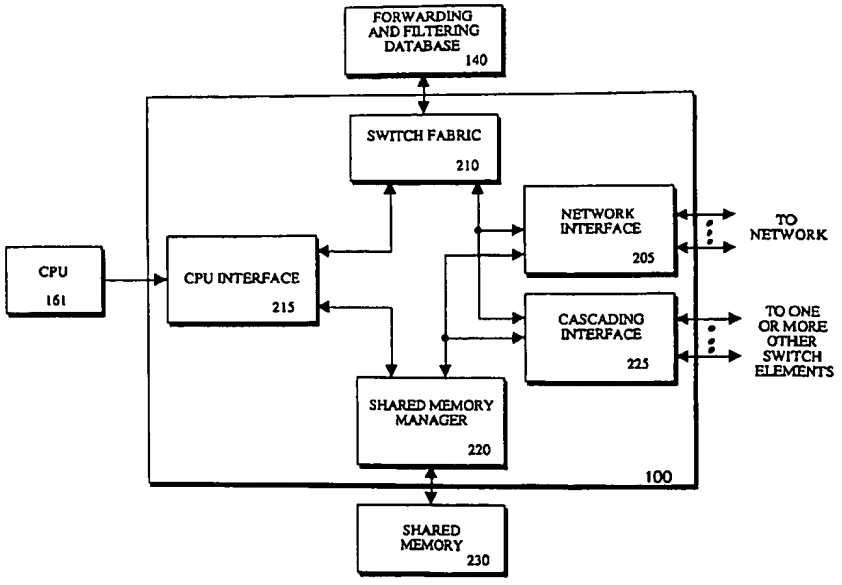

US 20020085560A1

(19) **United States**

(12) **Patent Application Publication** (10) **Pub. No.: US 2002/0085560 A1**

Cathey et al. (43) **Pub. Date: Jul. 4, 2002**

packet classificatio

(54) **PROGRAMMABLE PACKET PROCESSOR WITH FLOW RESOLUTION LOGIC**

(76) Inventors: **Jim Cathey**, Greenacres, WA (US); **Timothy S. Michels**, Greenacres, WA (US)

Correspondence Address:
CHRISTIE, PARKER & HALE, LLP
350 WEST COLORADO BOULEVARD
SUITE 500
PASADENA, CA 91105 (US)

(21) Appl. No.: **09/751,194**

(22) Filed: **Dec. 28, 2000**

Related U.S. Application Data

(63) Non-provisional of provisional application No. 60/206,617, filed on May 24, 2000. Non-provisional of provisional application No. 60/206,996, filed on May 24, 2000. Non-provisional of provisional application No. 60/220,335, filed on Jul. 24, 2000.

Publication Classification

(51) **Int. Cl.**[7] **H04L 12/28**; H04L 12/56

(52) **U.S. Cl.** **370/392**; 370/465

(57) **ABSTRACT**

A programmable packet switching controller has a packet buffer, a pattern-match module, a programmable packet classification engine and an application engine. The packet buffer stores inbound packets, and includes a header data extractor to extract header data from the inbound packets and to store the extracted header data in a header data cache. The header data extractor also generates a header data cache index and provides it to the packet classification engine for it to retrieve the extracted header data. The packet classification engine has a decision tree-based classification logic for classifying a packet. Each of the leaves of the tree represents a packet classification. The packet classification engine uses the header data cache index to retrieve the header data to perform multiple header checks, starting at a root of the tree and traversing branches until a leaf has been reached. The application engine has a number of programmable sub-engines arrayed in a pipelined architecture. The packet classification engine provides start indicators based on the packet classification to the programmable sub-engines to identify application programs to be executed. The sub-engines includes a source lookup engine, a destination lookup engine and a disposition engine, which are used to make a disposition decision for the inbound packets in a processing pipeline. The pattern match module is used to compare the packet to a pre-defined pattern in order to provide a disposition recommendation.

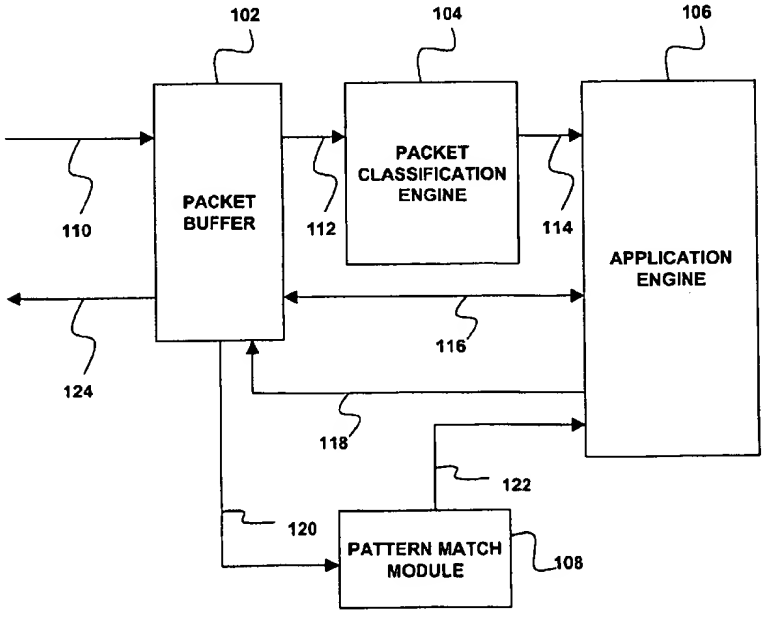

US 20020089929A1

(19) United States

(12) Patent Application Publication (10) Pub. No.: US 2002/0089929 A1

Tallegas et al. (43) Pub. Date: Jul. 11, 2002

(54) PACKET PROCESSOR WITH MULTI-LEVEL POLICING LOGIC

(76) Inventors: Mathieu Tallegas, Spokane, WA (US); Kelly Fromm, Verdale, WA (US); Dennis Paul, Liberty Lake, WA (US)

Correspondence Address:
CHRISTIE, PARKER & HALE, LLP
350 WEST COLORADO BOULEVARD
SUITE 500
PASADENA, CA 91105 (US)

(21) Appl. No.: 09/757,361

(22) Filed: Jan. 8, 2001

Related U.S. Application Data

(63) Non-provisional of provisional application No. 60/206,617, filed on May 24, 2000. Non-provisional of provisional application No. 60/206,996, filed on May 24, 2000. Non-provisional of provisional application No. 60/223,335, filed on Aug. 7, 2000.

Publication Classification

(51) Int. Cl.[7] H04J 3/14; H04L 12/56

(52) U.S. Cl. 370/230; 370/395.21

(57) ABSTRACT

A switch includes a backplane and multiple packet processors. One or more packet processors include multi-level policing logic. The packet processor receives a packet and classifies the packet into multiple policeable groups. The packet is compared against bandwidth contracts defined for the policeable groups. Nested lookups are performed for the packet in a policing database to identify the multiple groups and to retrieve policing data for the multiple policeable groups. The policing results, which may be combined into a single policing result by taking the worst case policing result, are applied to disposition logic as recommendations, and are combined with other disposition recommendations to make a disposition decision for the packet.

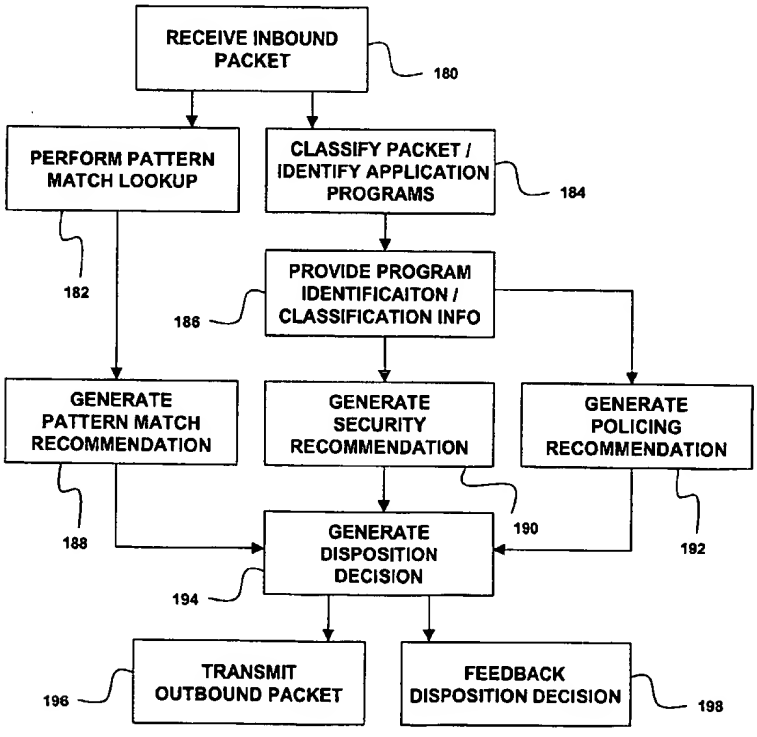

US 20030123448A1

(19) **United States**

(12) **Patent Application Publication** (10) **Pub. No.: US 2003/0123448 A1**

CHANG (43) **Pub. Date: Jul. 3, 2003**

(54) **SYSTEM AND METHOD FOR PERFORMING CUT-THROUGH FORWARDING IN AN ATM NETWORK SUPPORTING LAN EMULATION**

(76) Inventor: **CHI-HUA CHANG**, MILPITAS, CA (US)

Correspondence Address:
BLAKELY SOKOLOFF TAYLOR & ZAFMAN
12400 WILSHIRE BOULEVARD, SEVENTH FLOOR
LOS ANGELES, CA 90025 (US)

(*) Notice: This is a publication of a continued prosecution application (CPA) filed under 37 CFR 1.53(d).

(21) Appl. No.: **09/344,608**

(22) Filed: **Jun. 25, 1999**

Related U.S. Application Data

(60) Provisional application No. 60/090,939, filed on Jun. 27, 1998.

Publication Classification

(51) **Int. Cl.**[7] **H04L 12/28**; H04L 12/56

(52) **U.S. Cl.** **370/395.1**

(57) **ABSTRACT**

System and method in a network processor for performing cut-through forwarding of LANE packets without incurring the overhead associated with LANE protocol stack assisted routing. A content addressable memory (CAM) stores LEC uplink information including mapping between MAC destination addresses and VCC information. The network processor also stores LEC information table for corresponding VLAN identifiers and LECs. The LEC information table includes LEC ID information for the VLAN ID. For a LANE packet received from Ethernet and outbound to ATM destination, the network processor determines the LEC ID for the packet and then performs a CAM lookup to determine VCC information for the packet. The packet is then forwarded using the LEC ID and VCC information. For a LANE packet received from an ATM source and outbound to the Ethernet, the network processor determines the interface address of the packet. Based on the interface address, the network processor determines if the packet is a echoed or loopback packet or if its destination is the network processor itself. The packet is then routed to via the Ethernet if it is not a loopback or echo packet and the destination is not the network processor.

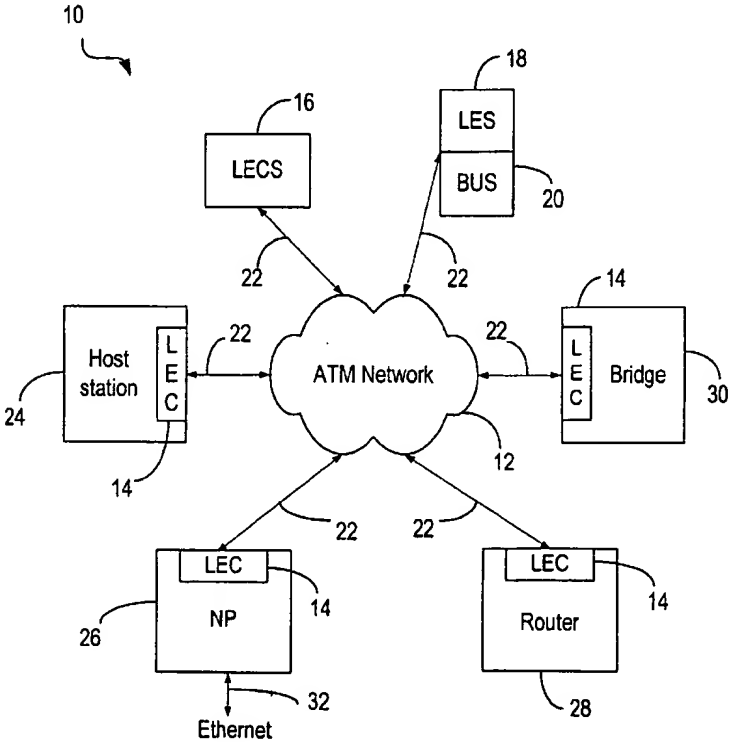

US 20020091795A1

(19) United States

(12) Patent Application Publication (10) Pub. No.: US 2002/0091795 A1

Yip (43) Pub. Date: Jul. 11, 2002

(54) METHOD AND SYSTEM OF AGGREGATE MULTIPLE VLANS IN A METROPOLITAN AREA NETWORK

(76) Inventor: Michael Yip, Sunnyvale, CA (US)

Correspondence Address:
BLAKELY SOKOLOFF TAYLOR & ZAFMAN
12400 WILSHIRE BOULEVARD, SEVENTH FLOOR
LOS ANGELES, CA 90025 (US)

(21) Appl. No.: 09/755,498

(22) Filed: Jan. 5, 2001

Publication Classification

(51) Int. Cl.[7] G06F 15/16

(52) U.S. Cl. 709/218; 709/201

(57) ABSTRACT

A method and system is provided in which data packets from multiple customer VLANs are forwarded over a MAN using VLAN aggregation. A layer-2 switch located at the edge of the MAN connects the customer VLANs to the MAN. The edge switch aggregates multiple customer VLANs (the "sub-VLANs") into one provider VLAN (the "super-VLAN"). When a packet is forwarded from the sub-VLAN to the super-VLAN and vice versa, the edge switch uses modified bridge forwarding rules to exchange the customer-configured VLAN-IDs with the provider-configured VLAN-IDs before transporting the packet over the MAN. The edge switch further uses modified bridge media access control (MAC) address learning rules to isolate one customer's traffic from another's (i.e. isolate one sub-VLAN's traffic from another sub-VLAN's traffic).

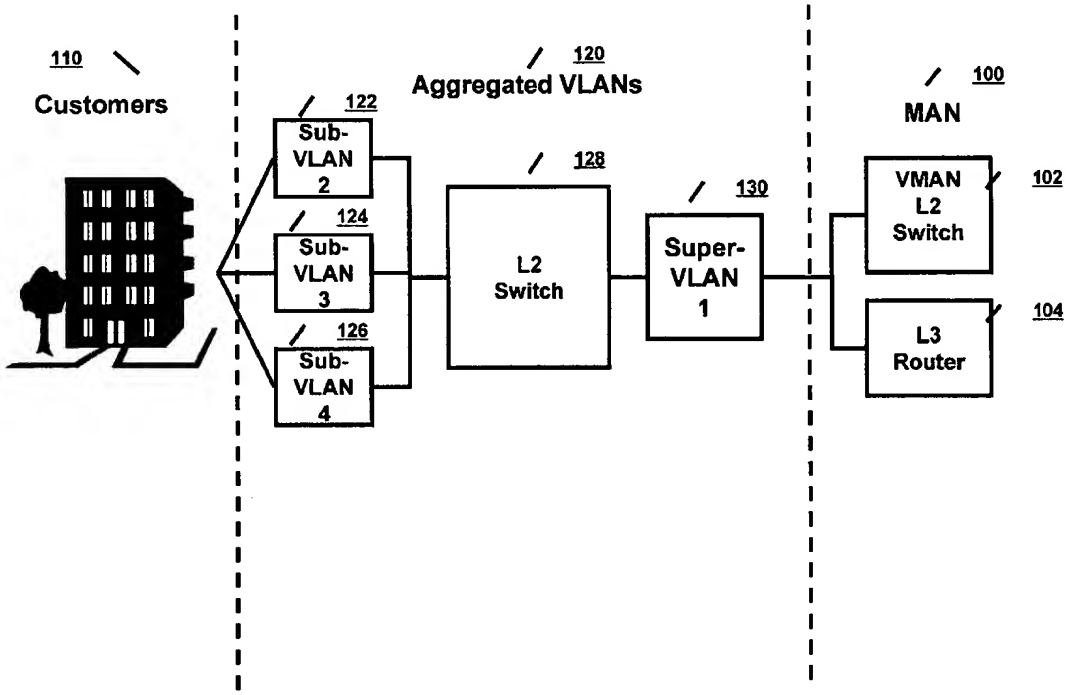

US006188694B1

(12) **United States Patent**
Fine et al.

(10) **Patent No.: US 6,188,694 B1**
(45) **Date of Patent: Feb. 13, 2001**

(54) **SHARED SPANNING TREE PROTOCOL**

(75) Inventors: **Michael Fine**, San Francisco, CA (US); **Silvano Gai**, Vigliano d'Asti (IT); **Keith McCloghrie**, San Jose, CA (US)

(73) Assignee: **Cisco Technology, Inc.**, San Jose, CA (US)

(*) Notice: Under 35 U.S.C. 154(b), the term of this patent shall be extended for 0 days.

(21) Appl. No.: **08/997,297**

(22) Filed: **Dec. 23, 1997**

(51) **Int. Cl.**[7] **H04L 12/28**
(52) **U.S. Cl.** **370/402**
(58) **Field of Search** 370/400, 401, 370/402, 403, 404, 357, 254, 255, 256

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,922,486		5/1990	Lidinsky et al.	370/60
5,018,137	*	5/1991	Backes et al.	370/401
5,394,402	*	2/1995	Ross	370/402
5,430,728	*	7/1995	Narayanan et al.	370/404
5,740,171		4/1998	Mazzola et al.	370/392
5,742,604		4/1998	Edsall et al.	370/401
5,761,435	*	6/1998	Fukuda et al.	709/238
5,764,636		6/1998	Edsall	370/401
5,790,808	*	8/1998	Seaman	709/223
5,818,838	*	10/1998	Backes et al.	370/390
5,878,232	*	3/1999	Marimuthu	709/249
5,956,335	*	9/1999	Backes et al.	370/392
6,032,194	*	2/2000	Gai et al.	709/239

OTHER PUBLICATIONS

Radia Perlman, "Interconnections: Bridges and Routers", pp. 54–64, Copr. 1992.*
S. Horowitz, "*Dual-Layer Spanning Tree*", (A Spanning Tree Proposal for IEEE 802.1Q), May 14, 1997.
Lucent Technologies Inc.'s Initial Disclosure of Prior Art Under Civil Local Rule 16–7, Case No. C98–20836 JW (PVT) ENE.
M. Seaman and D. Delaney,*Single or Multiple Filtering Databases* (May 8, 1997).
Radia Perlman, "*Interconnections: Bridges and Routers*", pp. 54–64, Copr. 1992.
IEEE, "*Draft Standard for Virtual Bridged Local Area Networks*", pp. 6–7 and 28–31 (Jul. 1996).
IEEE, "*Draft Standard for Virtual Bridged Local Area Networks*", pp. 1–7 and 35–39 (Aug. 1996).
IEEE, "*Draft Standard for Virtual Bridged Local Area Networks*", pp. 1–10 and 20–21 (May 1997).
Steven Horowitz, IEEE, "*Dancing Bears*", IEEE Working Draft (Mar. 1996).

* cited by examiner

Primary Examiner—Melvin Marcelo
Assistant Examiner—Phuongchau Ba Nguyen
(74) *Attorney, Agent, or Firm*—Cesari and McKenna, LLP

(57) **ABSTRACT**

A shared spanning tree protocol (SSTP) creates a plurality of spanning trees (i.e., loop-free paths) which are shared among one or more virtual local area network (VLAN) designations for data transmission within a computer network. Each shared spanning tree includes and is defined by a primary VLAN and may be associated with one or more secondary VLANs. In order to associate VLAN designation(s) with a single shared spanning tree, network devices exchange novel shared spanning tree protocol data units (SST-PDUs). Each SST-PDU corresponds to a given primary VLAN and preferably includes one or more fields which list the secondary VLAN designations associated with the given primary VLAN. The association of VLAN designations to shared spanning trees, moreover, preferably depends on which path traffic is to follow as well as the anticipated load characteristics of the various VLANs. The association of VLAN designations to shared spanning trees thus provides a degree of load balancing within the network. Data messages tagged with a particular VLAN designation are then distributed by the devices only along the shared spanning tree to which that VLAN has been associated by SSTP.

38 Claims, 7 Drawing Sheets

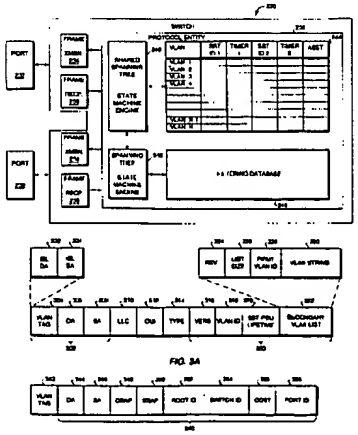

US006445709B1

(12) **United States Patent**
Chiang

(10) **Patent No.: US 6,445,709 B1**
(45) **Date of Patent: Sep. 3, 2002**

(54) **METHOD AND APPARATUS FOR FINDING A MATCH ENTRY USING RECEIVE PORT NUMBER EMBEDDED IN THE PORT VECTOR**

(75) Inventor: **John M. Chiang**, San Jose, CA (US)

(73) Assignee: **Advanced Micro Devices, Inc.**, Sunnyvale, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **09/311,367**

(22) Filed: **May 13, 1999**

(51) **Int. Cl.**[7] **H04L 12/56**; G06F 13/00

(52) **U.S. Cl.** **370/399**; 709/236; 706/12

(58) **Field of Search** 709/236; 370/397, 370/399; 706/12, 25

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,852,607	A *	12/1998	Chin	370/401
5,892,912	A *	4/1999	Suzuki et al.	370/395.53
6,104,696	A *	8/2000	Kadambi et al.	370/218
6,111,874	A *	8/2000	Kerstein	370/389
6,128,296	A *	10/2000	Daruwalla et al.	370/389
6,151,297	A *	11/2000	Congdon et al.	370/216
6,154,446	A *	11/2000	Kadambi et al.	370/239
6,161,144	A *	12/2000	Michels et al.	709/238
6,236,654	B1 *	5/2001	Egbert	370/392
6,243,394	B1 *	6/2001	Deng	370/466
6,266,705	B1 *	7/2001	Ullum et al.	709/238
6,335,935	B2 *	1/2002	Kadambi et al.	370/396
6,356,551	B1 *	3/2002	Egbert	370/389
2001/0012294	A1 *	8/2001	Kadambi et al.	370/391
2001/0043611	A1 *	11/2001	Kadambi et al.	370/429

* cited by examiner

Primary Examiner—Kenneth R. Coulter

(57) **ABSTRACT**

A network switch configured for switching data packets across multiple ports uses an address table to generated frame forwarding information. The switch receives frame information including a source address and destination address along with a virtual local area network (VLAN) ID, if applicable. A decision-making engine searches a network address table to "learn" Medium Access Control (MAC) addresses without having to store receive port numbers by embedding the receive port number in the stored port vector fields.

15 Claims, 12 Drawing Sheets